

Performance Impacts of Node Failures on a Chord-Based Hierarchical Peer-to-Peer Network

Quirin Hofstätter

Lehrstuhl für Kommunikationsnetze
Fakultät für Elektrotechnik und Informationstechnik
Technische Universität München, Germany
`quirin.hofstaetter@tum.de`

Abstract. Peer-to-Peer networks are designed to provide decentralized, fault-tolerant alternatives to traditional client-server architectures. In previous work, the resilience of structured P2P networks has been evaluated analytically as well as by simulation. In this paper, we discuss the influence of peer failures on a core component of our *cost-optimized* hierarchical P2P protocol *Chordella*. We show that the used algorithms lead to a working system even if a high number of superpeers is subject to failures.

1 Introduction

Peer-to-Peer (P2P) networks are based on the construction of an overlay network on the application layer, where the decentralized protocol allows the users to share certain resources. Current generations of P2P networks are based on Distributed Hash Tables (DHTs), used to maintain a consistent, global view of the available resources. Popular protocols are e.g. Chord [1] or Kademlia [2]. However, the stability of these overlay networks is directly affected by user behavior (joining or leaving the network) and the performance of the network connection. When many peers fail (e.g. mobile peers losing power supply or the wireless network link), the P2P network may be split into a number of disjoint networks or even break down completely. Hence a careful design of the used algorithms is crucial.

A number of previous works on the stability of P2P networks has focused the Chord protocol and its behavior in non-optimal conditions [3,4]. The authors examine the behavior of the overlay in settings with inconsistent routing table entries, the recovery mechanisms and their boundaries.

In our previous work [5,6], we extended standard Chord to a two-tier hierarchical P2P protocol for mobile use: *Chordella*. High performance nodes act as superpeers (SP) forming the DHT, nodes with less power (e.g. mobile phones) are attached to these superpeers as leaf nodes (LN). In order to drive the system in a cost-optimal state (in terms of network traffic), we introduce metrics and algorithms to determine, reach and maintain the optimal number of superpeers in the network in a fully distributed way.

2 Impact on Chordella’s Algorithms

In [6] we describe the used algorithm to achieve and maintain the cost-optimal ratio between superpeers and leaf nodes in the network. We show that it is able to drive the system in a state that the relative deviation between the theoretical optimum and the value measured in simulation does not exceed 5% in a realistic churn scenario. In the worst case it stays within 10%. These scenarios include random failures of peers in the network. In this work, we study the impact of non-random failures on the introduced algorithms. We use the same parameters for simulation as in [6], due to space limitations we refer to that publication for details. Of course, the failure scenario differs: We let a differently sized number of peers fail every twelve stabilization periods $T_{stab} = 5$ s. The removed peers build a block in the ID space, i.e. they are direct successors. Hence the Chord overlay experiences a worst-case failure. All results are averages of ten independent simulation runs, confidence intervals are too small to be shown.

Figure 1 shows the number of superpeers against the simulation time in a case where no peers fail (best case, solid line) and with equally distributed failures (as in [6], dashed line). You can see the three phases of the simulation: the join phase (0 to 7200 seconds) where the nodes join the network at a constant rate, the churn phase where nodes join and leave the network (normal operation mode) and the leave phase where all nodes leave the network at a constant rate. In this scenario the system is able to maintain a number of superpeers very close to the optimum. This random failure case is the reference for the following scenarios.




Fig. 1. Number of SPs: No failures (solid) **Fig. 2.** Number of SPs: Random failures / Equally distributed failures (dashed) (solid) / 10% block (dashed)




Fig. 3. Number of SPs: Random failures (solid) / 25% block (dashed)




Fig. 4. Signaling traffic in different failure situations: none/random/block

Figure 2 shows the case where in the operational phase (churn phase) a successional block of 10% of all superpeers is selected every $12 \cdot T_{\text{stab}}$ and removed from the network compared to the random failure case. This leads to an immediate drop in the total number of superpeers. Although some overcompensations (due to the distributed nature of the used algorithm) are observable, the algorithm is able to compensate this error and regain a stable state within the twelve stabilization periods.

When the size of the block is increased to 25% of the superpeers as shown in Figure 3, a similar result is obtained. Both, the drop and the overcompensation spikes are larger but again the system manages to return to a stable state.

Since the goal of the algorithms is a cost-optimal operation, it makes sense to also evaluate the costs (i.e. the signaling traffic) in the different failure scenarios. Figure 4 shows the comparison of signaling messages per second in the network. Obviously the introduction of failures leads more signaling traffic. Uniformly distributed failures only introduce a moderate increase compared to the best case without any disturbance. The introduced block-failures lead to 125% more signaling messages compared to the random failures if 10% of the nodes fail or even 670% in the 25%-block case.

3 Conclusion

The introduction of worst-case failures in the superpeer tier of the Chordella system leads to deviations from the optimal number of superpeers in the system. However, our algorithm is able to compensate these interferences within twelve stabilization periods. The number of messages needed to maintain a working system increases rapidly with the ratio of failed peers. Of course, the failures have also an impact on parameters like the query success and query duration. These results are omitted due to space constraints.

References

1. Stoica, I., Morris, R., Karger, D., Kaashoek, M., Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. In: Proceedings of the 2001 ACM SIGCOMM, p. 160. ACM, New York (2001)
2. Maymounkov, P., Mazieres, D.: Kademlia: A peer-to-peer information system based on the xor metric. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, p. 2. Springer, Heidelberg (2002)
3. Leonard, D., Rai, V., Loguinov, D.: On lifetime-based node failure and stochastic resilience of decentralized peer-to-peer networks. In: Proceedings of the 2005 ACM SIGMETRICS, pp. 26–37. ACM, New York (2005)
4. Binzenhöfer, A., Staehle, D., Henjes, R.: On the Stability of Chord-based P2P Systems. University of Würzburg 347 (2004)
5. Zoels, S., Despotovic, Z., Kellerer, W.: On hierarchical DHT systems—An analytical approach for optimal designs. *Computer Communications* 31(3), 576–590 (2008)
6. Zoels, S., Hofstätter, Q., Despotovic, Z., Kellerer, W.: Achieving and Maintaining Cost-Optimal Operation of a Hierarchical DHT System. In: Proceedings of the 2009 IEEE International Conference on Communications (ICC) - Next Generation Networking Symposium (2009)